ETH zürich

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

# A Look Into PULP Publications

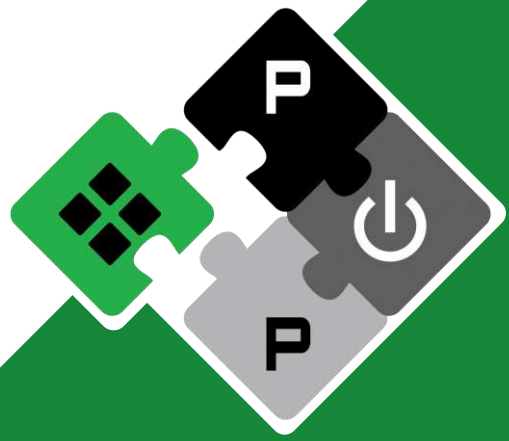

**Davide Rossi** davide.rossi@unibo.it

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform

**PULP Platform**

Open Source Hardware, the way it should be!

# PULP: The Origins

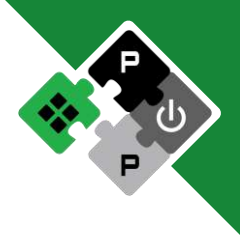


- **Project started in 2013**
- **A collaboration between University of Bologna and ETH Zürich**
  - We wanted to start with a clean slate, no need to remain compatible to legacy systems, no dependencies with any commercial IP
- **Original Academic/Research goal**
  - Push energy efficiency of IoT computing systems as much as possible
  - Create a compute platform and an ecosystem used for research on Computer Architectures by our team as well as other groups in the World
- **Original Approach**
  - Exploit Open Source Ips as much as possible
  - We'll never do a processor
  - We'll never do a compiler

# Phase 0: Open (And Close) Source Processors Landscape

- **Reasonably Good Quality IPs**
- **Area around 50 kGates**
- **Open Source**
- **Compiler Support**
- **Community support**

- **We decided to go with OR1200**
- **Issues:**
  - No single cycle data interface
  - No Open Source ISS

Architectural Requirements:

- DTCM interface
- PCACHE
- Blocking bus interface
- User-defined extension interface
- Floating point unit
- NOT obfuscated RTL
- Instruction Set Simulator

A12000

| PROCESSOR | PCACHE | BUS | EXT IF | F.P. UNIT | RTL VISIBLE | ISS | AREA (Kg) |
|---|---|---|---|---|---|---|---|
| STxp70 | YES | T3 | YES | YES | YES | YES | 42+ |
| REISC4 | YES | CUSTOM | NO | NO | ? | ? | ? |
| ARM CORTEX M0+ | NO | AMBA | NO | NO | (NO) | YES | ~9 |
| ARM CORTEX M3 | NO | AMBA | NO | NO | (NO) | YES | ~27 |
| ARM CORTEX M4 | NO | AMBA | NO | YES | (NO) | YES | ~38 (FPU) |
| ARC EM4 | NO | AMBA | YES | YES | (NO) | YES | 10 - 20 |
| ARC EM6 | YES | AMBA | YES | YES | (NO) | YES | 20 - 40 |
| LATTICE Mico32 | YES | WISHB. | NO | NO | NO | NO | ~30 |
| OPENRISC | YES | WISHB. | NO | YES | YES | NO | ~50 |

TARGET GCC LLVM → EMULATORS SOFTWARE

COMPLICATIONS BUONA QUALITÀ OPEN

# The Community Supporting OR1200 @ ORCONF 2013


ORCONF 2013 5-6/10/2013

# OR1200 Microarchitectural Issues

- Blocking LSU (2 cycles)
- Blocking Multiplier (3 cycles)
- I$ →TCDM Combinational Path
- No Support for compressed instructions

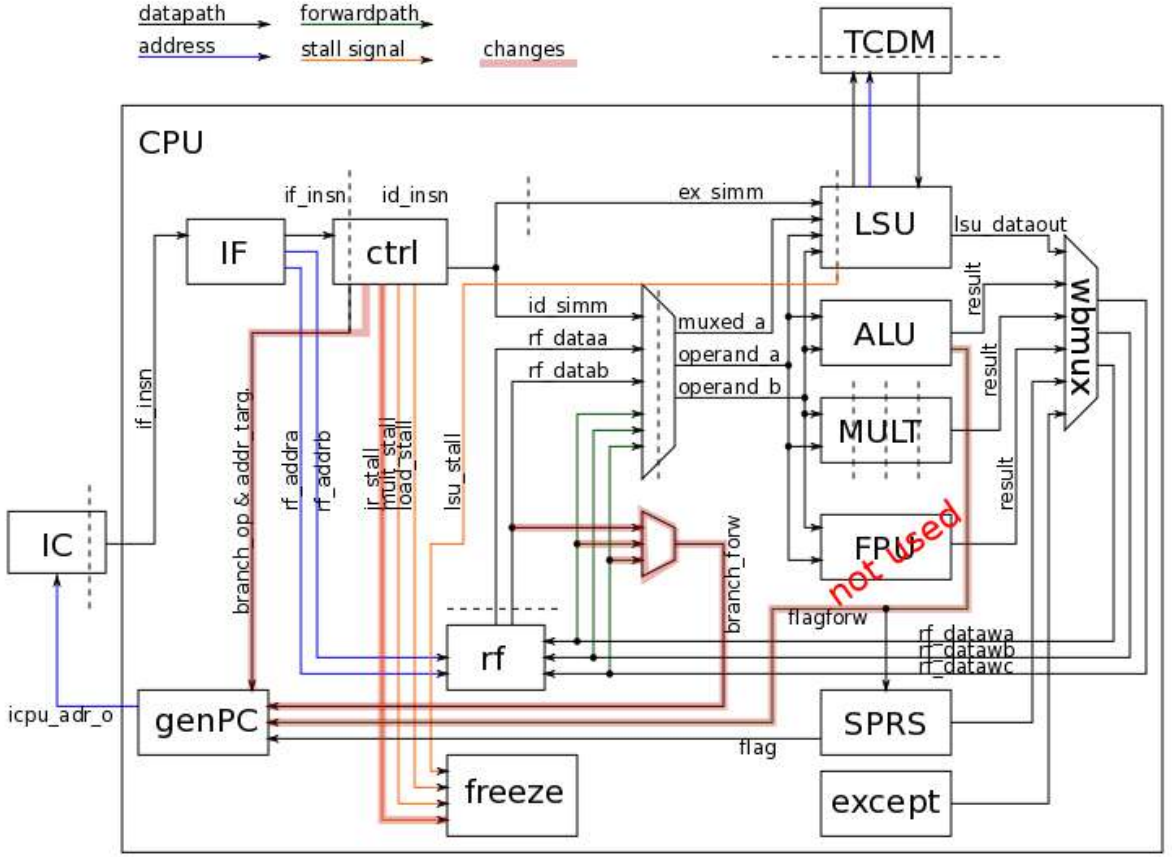

# First PULP Publication: Michael @ GLSVLSI

## Customizing an Open Source Processor to Fit in an Ultra-Low Power Cluster with a Shared L1 Memory

Michael Gautschi[1], Davide Rossi[2], Luca Benini[1,2]

[1] Integrated Systems Laboratory
ETH Zurich, Switzerland
{gautschi,benini}@iis.ee.ethz.ch

[2] Electrical, Electronic, and Information Engineering
University of Bologna, Italy
{davide.rossi, luca.benini}@unibo.it

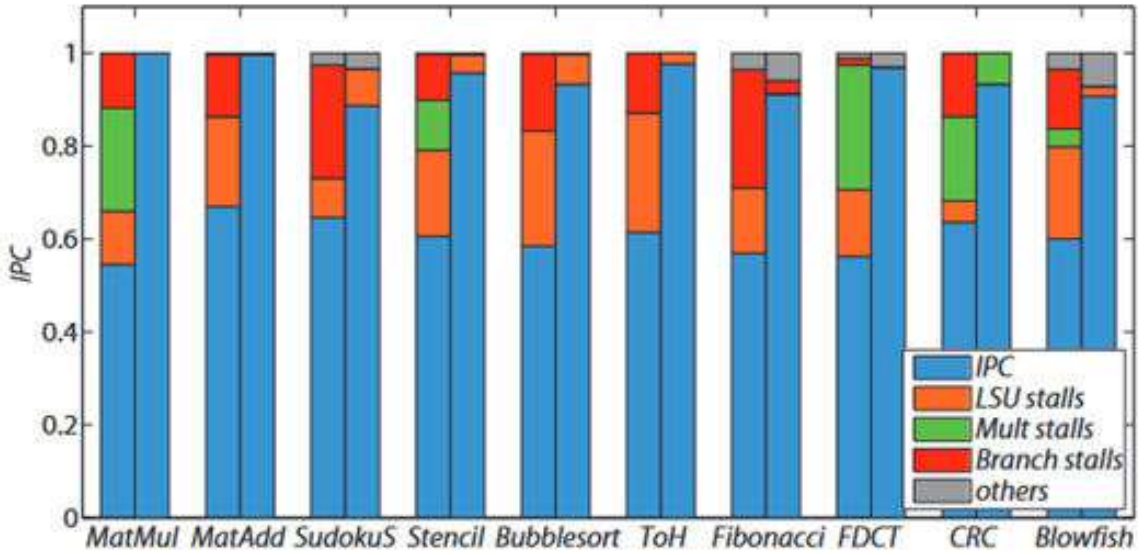


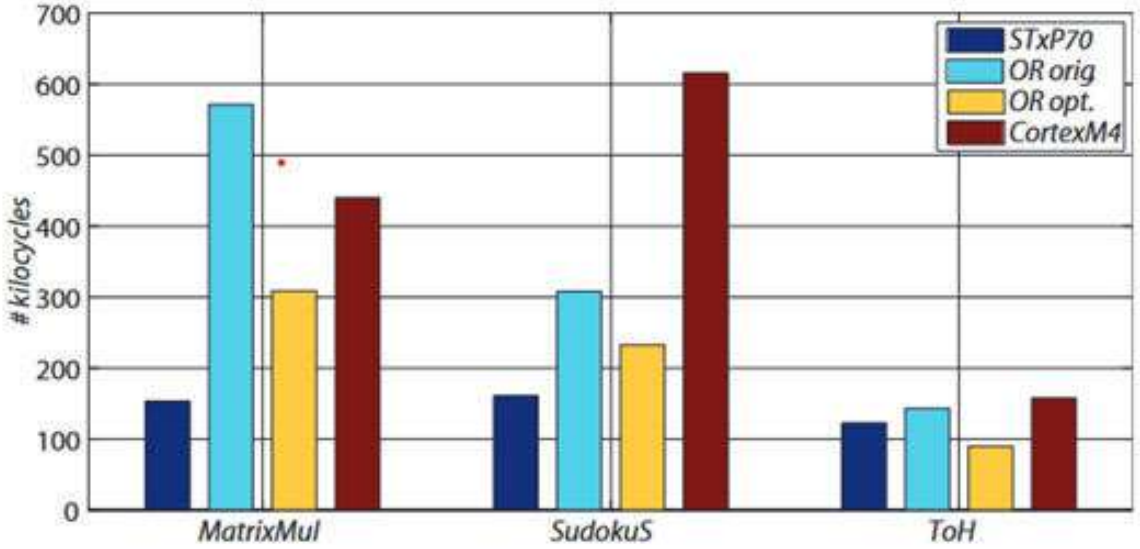

# The first “PULP cluster”: PULPv1 - TO Dec 2013

- 4-cores + I$ cluster
- Logarithmic interco.
- I$ + L2 Interco.
- L2 memory
- JTAG

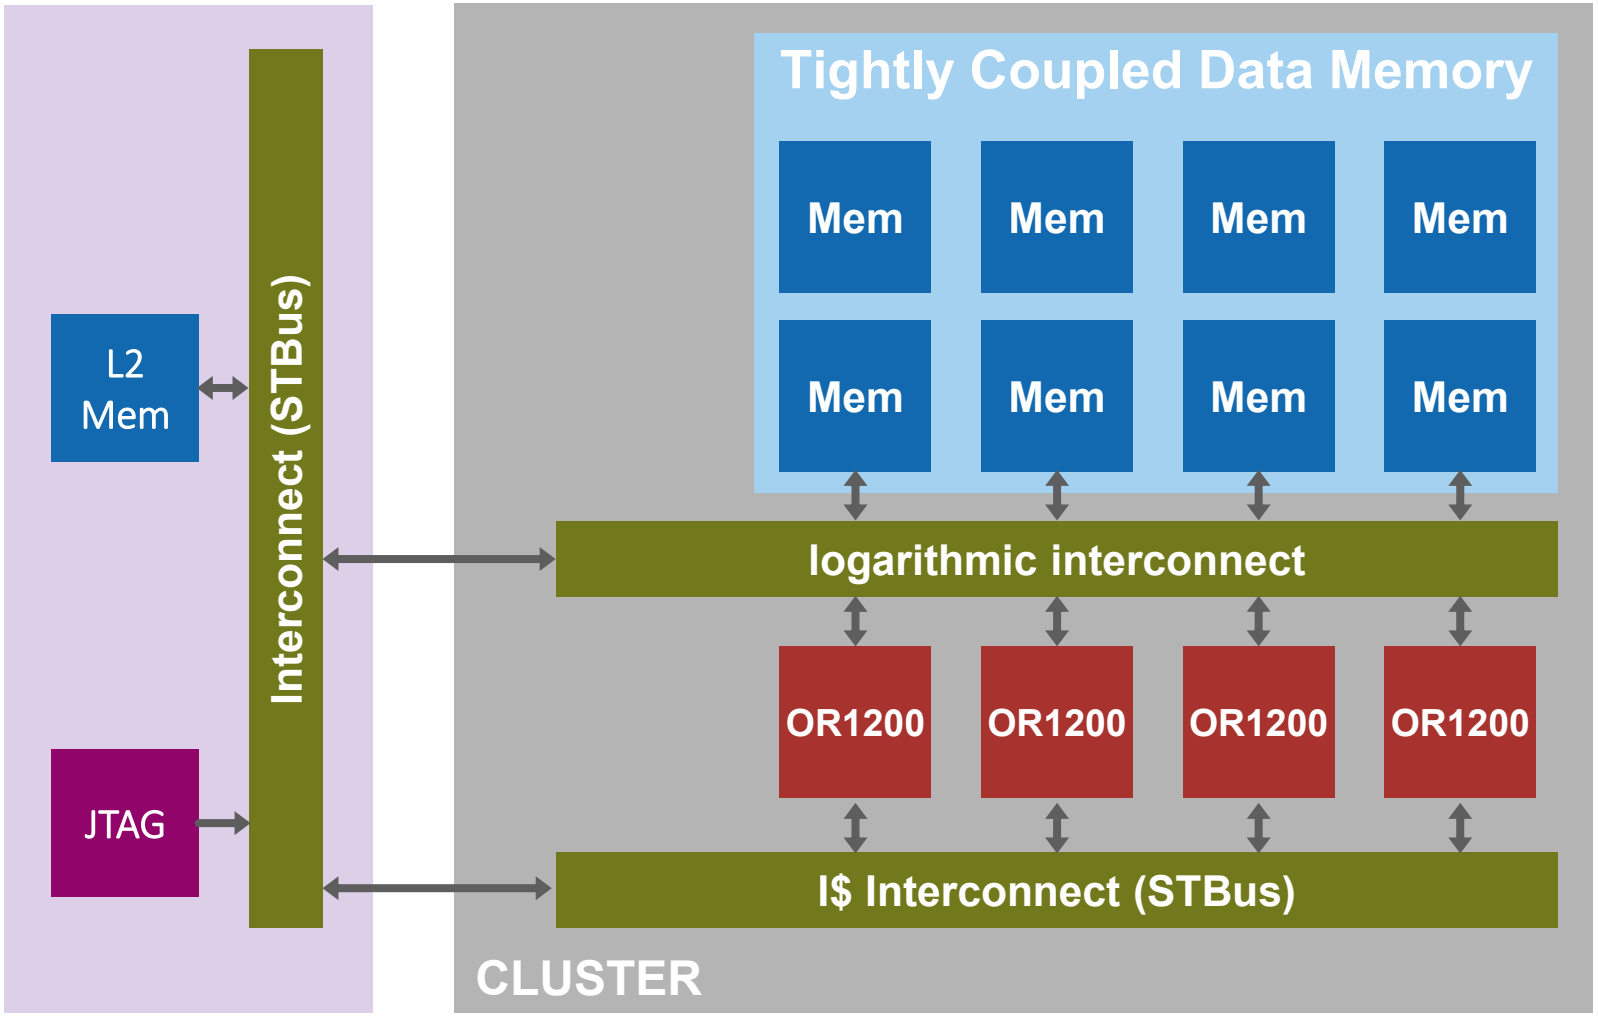


....not really “Open Source”...☺

# The Evolution of the Species: PULPv2 - TO Dec 2014

- PULPv1 +
- New (very rudimental) but Open Source AXI interconnect
- Some new peripherals (SPIM, SPIS)
- Latch based I$
- DVFS-ready

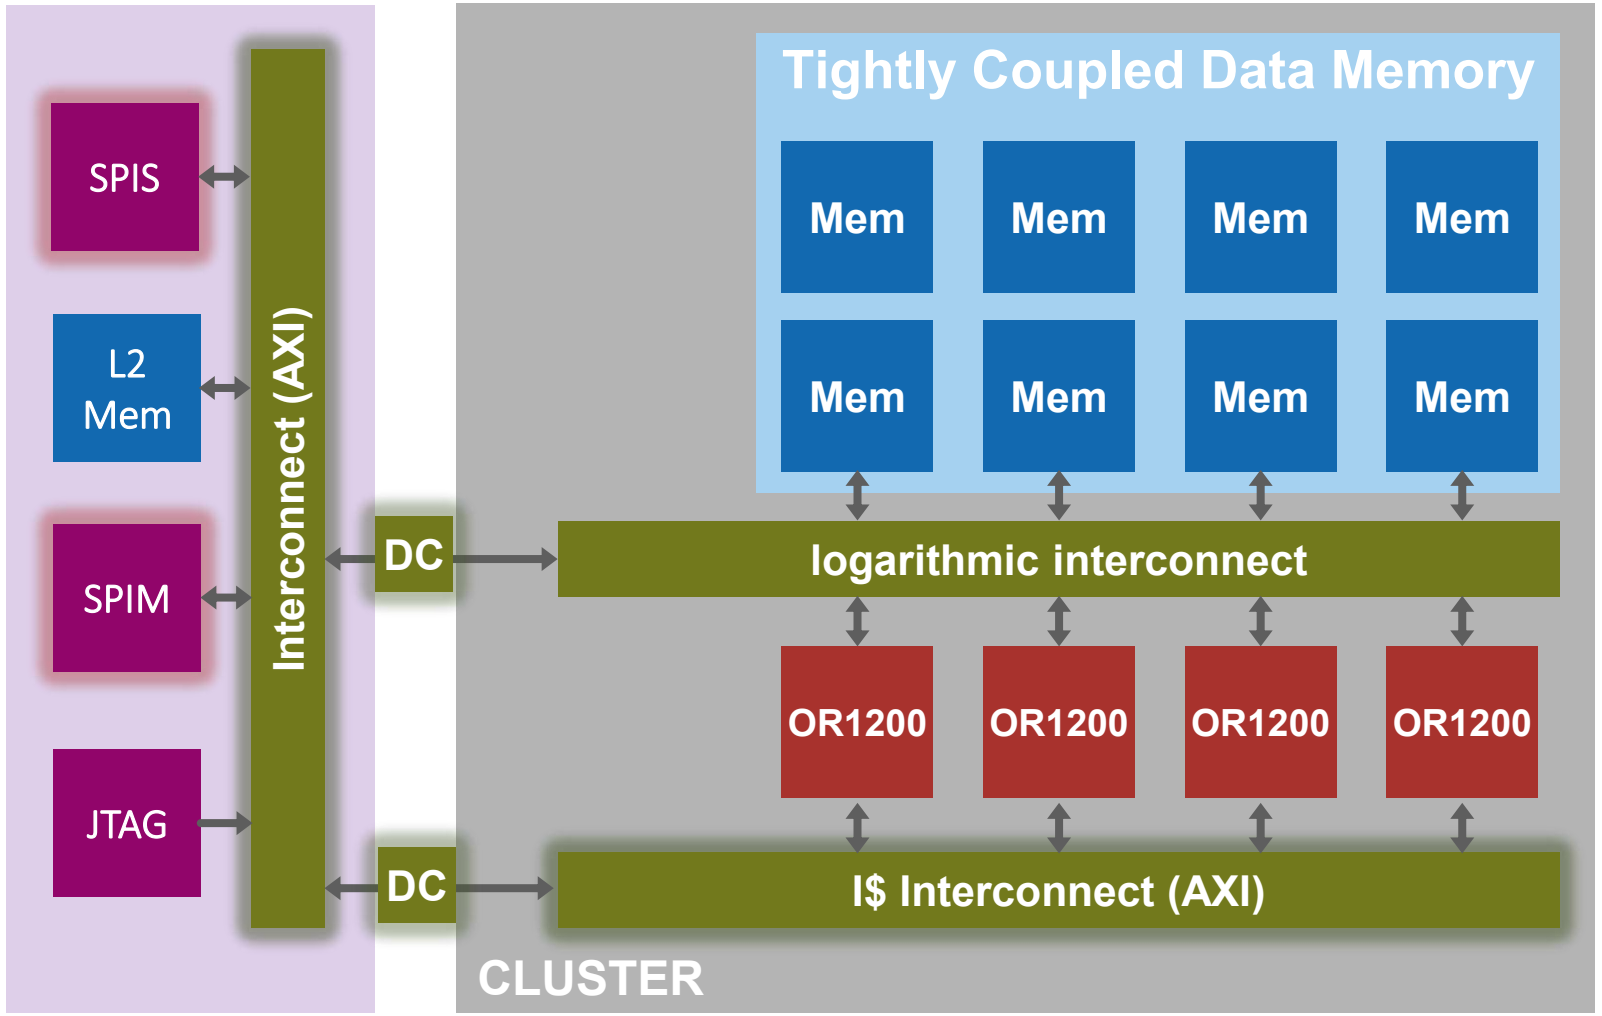

# The Evolution of the Species: PULPv3 - TO Dec 2015

- PULPv2 +
- OR1ON cores with DSP Extensions
- Hardware Synchronizer
- Shared Instruction Cache (latch based)
- I/O DMA
  - SPI M, SPI S, UART, I2C, I2S, CAM

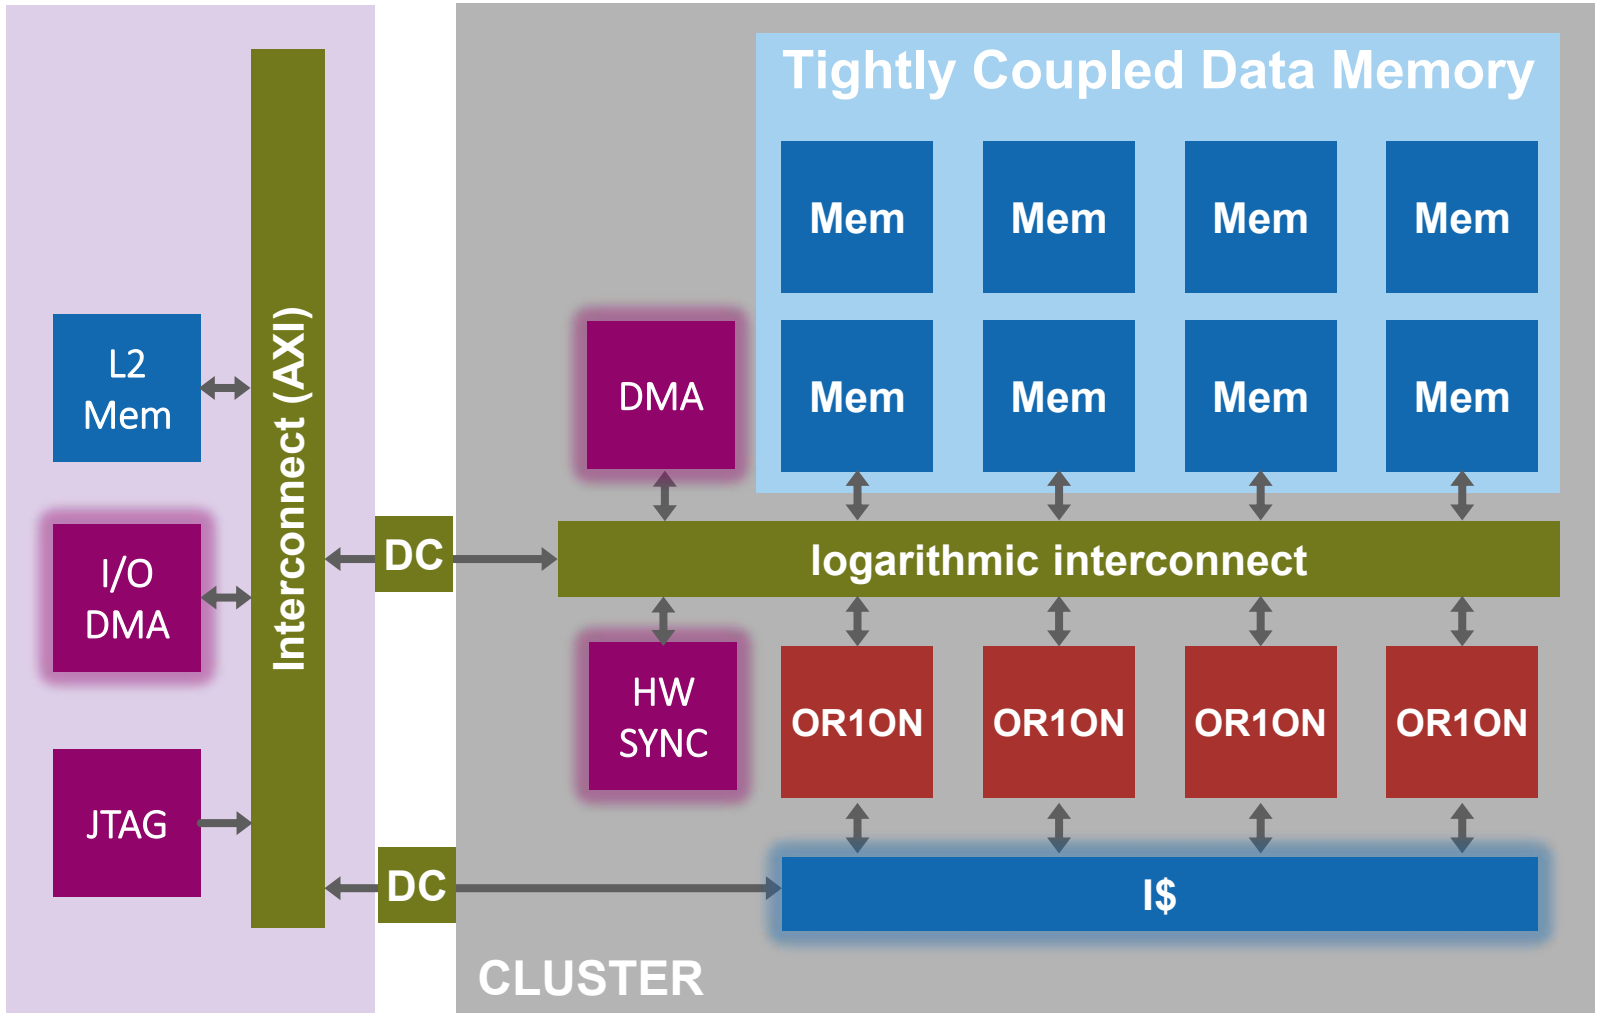


In the mean time…something was happening on the other side of the world…

# Wind of Change: First RISC-V Workshop, January 2015





January 14-15, 2015
Marriott Hotel, Monterey, CA

# First Big Conference (PULPv1 →v3) @ Hot Chips 27, Aug. ‘15

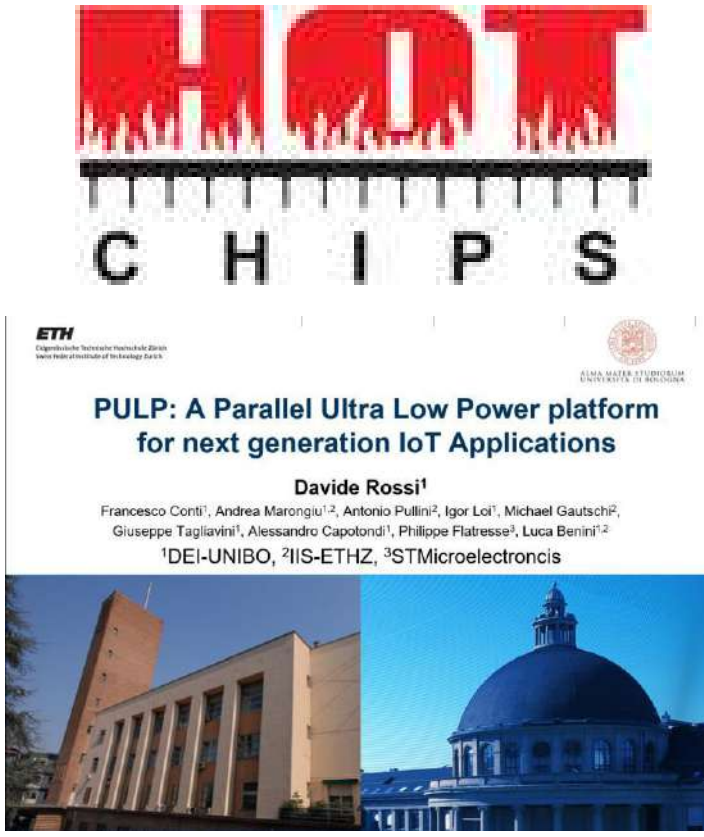


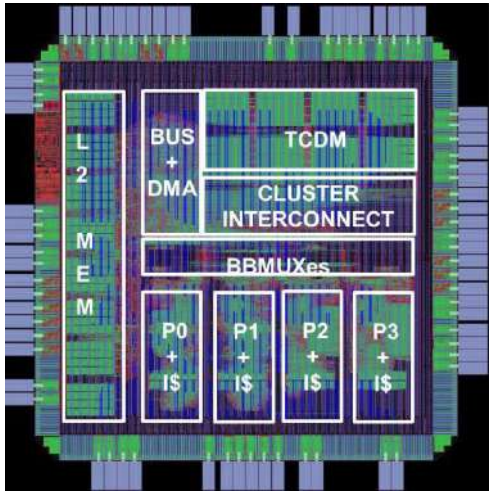


Tape out: Dec 2013

| CHIP FEATURES | |
|---|---|
| Technology | 28nm FDSOI (RVT) |
| Area | 3mm² |
| # Cores | 4xOpenRISC |
| I$ | 4x1kbyte (private) |
| TCDM | 16 kbyte |
| L2 | 16 kbyte |
| BB regions | 4 |
| VDD range | 0.45 – 1.2V |
| VBB range | -1.8V - +0.9V |

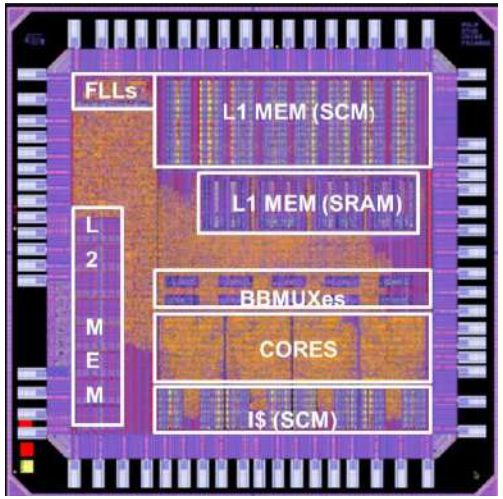


Tape out: Dec 2014

| CHIP FEATURES | |
|---|---|
| Technology | 28nm FDSOI (LVT) |
| Area | 3mm² |
| # Cores | 4xOpenRISC |
| I$ (SCM) | 4x1kbyte (private) |
| TCDM | 32 + 8 Kbyte |
| L2 | 64 kbyte |
| BB regions | 8 |
| VDD range | 0.3 - 1.2V* |
| VBB range | 0V - 2V* |
| GOPS | 4* |
| GOPS/W | 135* |

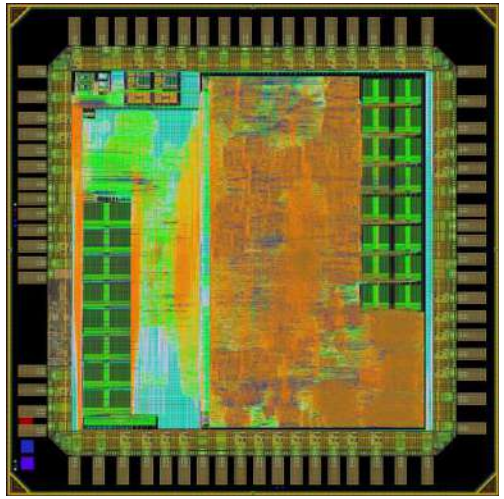

Tape out: Dec 2015

| CHIP FEATURES | |
|---|---|
| Technology | 28nm FDSOI (RVT) |
| Area | 3mm² |
| # Cores | 4xOR1ON |
| I$ | 4 kbyte (shared) |
| TCDM | 64 + 8 kbyte |
| L2 | 128 kbyte |
| BB regions | 2 |
| VDD range | 0.5 – 0.7V* |
| VBB range | -1.8V - +0.9V |
| GOPS | 2.2 |
| GOPS/W | 400 |

## Focus on Exploitation of FD-SOI technology (i.e. body bisasing)

ETH zürich ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

# Wind of Change...Open-Source Contributions at Hot Chips 27


# PULP Cluster with Shared LNU @ ISSCC 2016 (Feb '16)



## 4.6 A 65nm CMOS 6.4-to-29.2pJ/FLOP@0.8V Shared Logarithmic Floating Point Unit for Acceleration of Nonlinear Function Kernels in a Tightly Coupled Processor Cluster

Michael Gautschi[1], Michael Schaffner[1], Frank K. Gürkaynak[1], Luca Benini[1,2]

[1]ETH Zurich, Zurich, Switzerland, [2]University of Bologna, Bologna, Italy

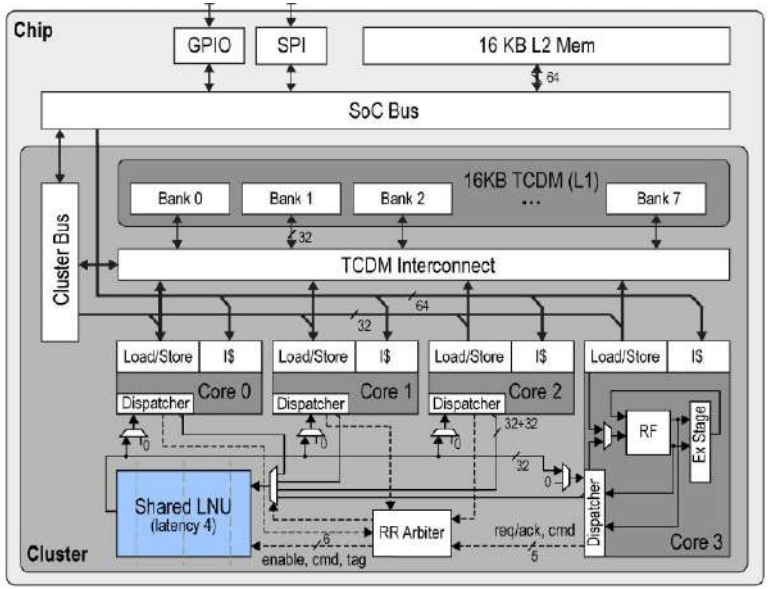


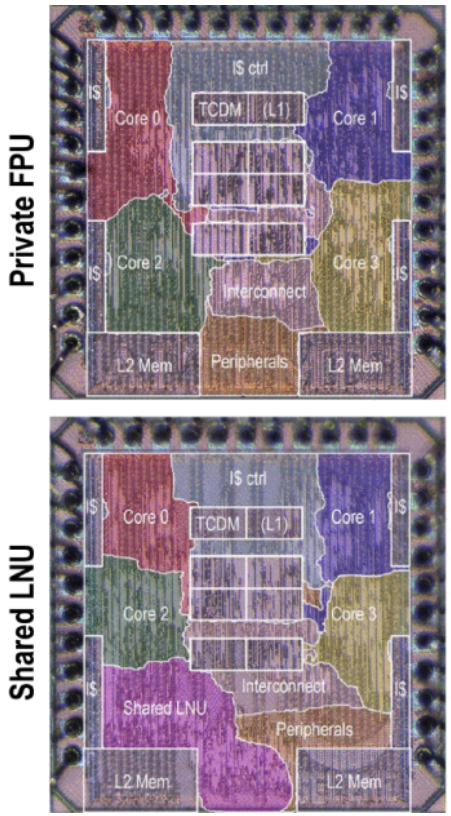


| Implementation Details | | Private FPU | Shared LNU | ELM [1,2] |
|---|---|---|---|---|
| Technology | | 65nm LVT | 65nm LVT | 180nm |
| max speed [MHz] | | 374 | 337 | 125 |
| max. Throughput [GFLOPS] | | 1.1 | 0.9 | 0.084 |
| Power @100MHz, 1.2V, 25°C [mW] | | 41.84 | 44.0 | - |
| Leakage @ 1.2V, 25°C [mW] | | 2.823 | 3.019 | - |
| Best FP-operator efficiency @0.8V [pJ/FLOP] | | 15.2 | **6.4** | - |
| Precision (max err) [ulp] | | 0.5 | 0.478[1] | 0.454[1] |
| avg. lnu/fpu utilization | | 0.21 | 0.37 | - |
| Total area [kGE] | | 719 | 749 | - |
| Single core area [kGE] | | 51.1[2] | 44.5 | - |
| **Instruction support** | | **Private FPU** | **Shared LNU** | **ELM [1,2]** |
| Latency add/sub/casts | hw | 2/2/2 | 4/4/4 | 3/3(4)/ - |
| Latency mul/div/sqrt [3] | hw | 2/-/- | 1/1/1 | 1/1/1 |
| | sw | -/62/56 | -/-/- | -/-/- |
| Latency exp/log [3] | hw | -/- | 4/4 | - / - |
| | sw | 51/85 | -/- | -/- |

[1] precision compared to FP [3] [2] including FPU [3] sw emulations: *div* (range reduction, linear initial estimate and 3 Newton iterations), *sqrt* (fast-inverse sqrt and 3 Newton iterations), *log* and *exp* (range reduction, polynomials from "Computer Approximations", J. Hart, 1978, ISBN 0882756427)

# Fulmine: Secure IoT Endpoint SoC (Jan 2016)

- Focused on Architecture
- 4 OR1ON cores
- L2: 192 kB
- TCDM: 64 kB
- HWPEs:
  - HWCE
  - HW CRYPT

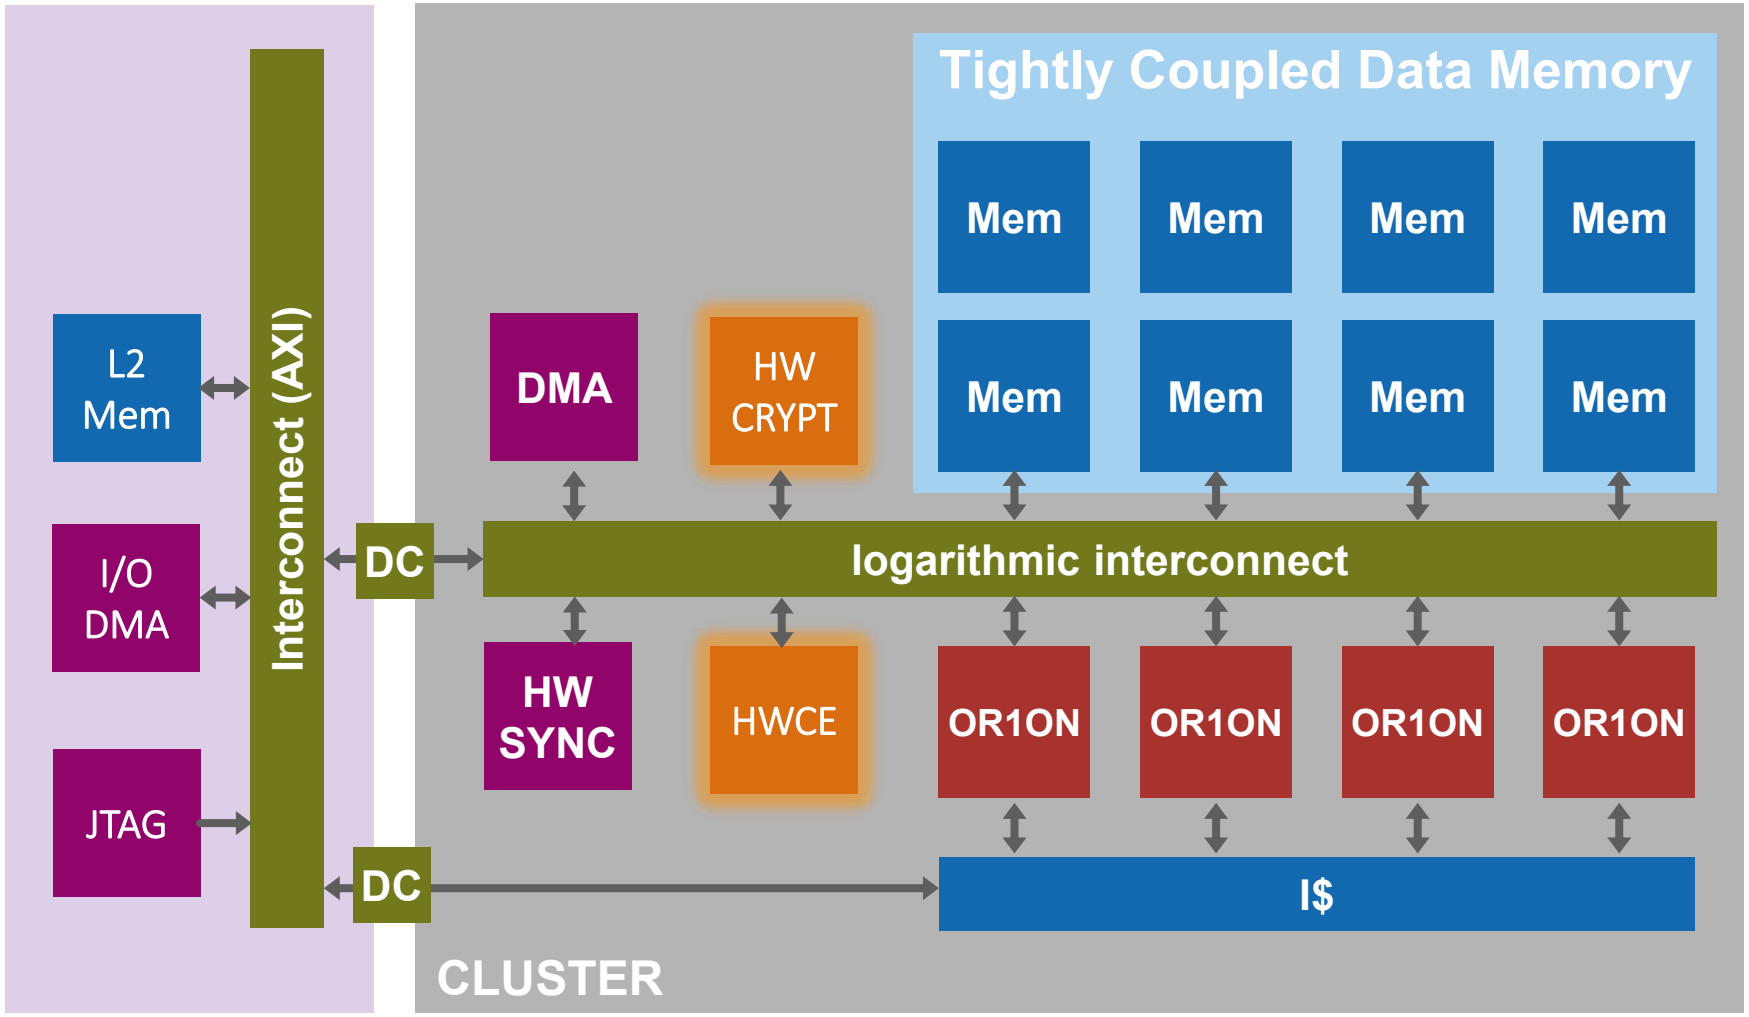

# Fulmine: Pushing Heterogeneity with HWPEs!

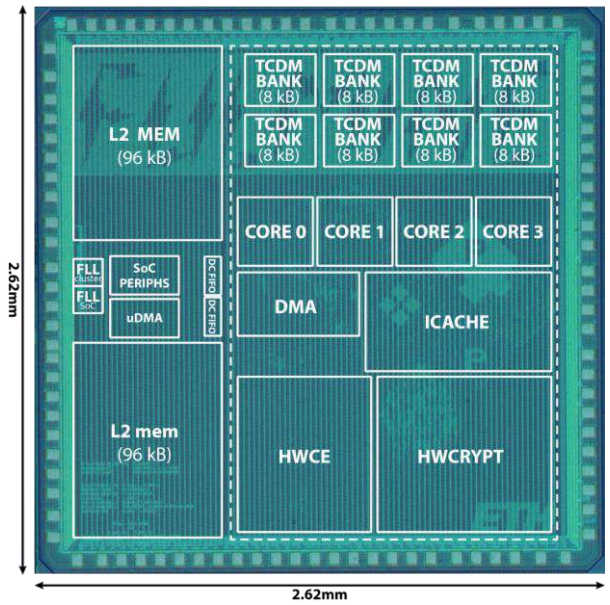


## Secure Autonomous Aerial Surveillance: 45x less energy

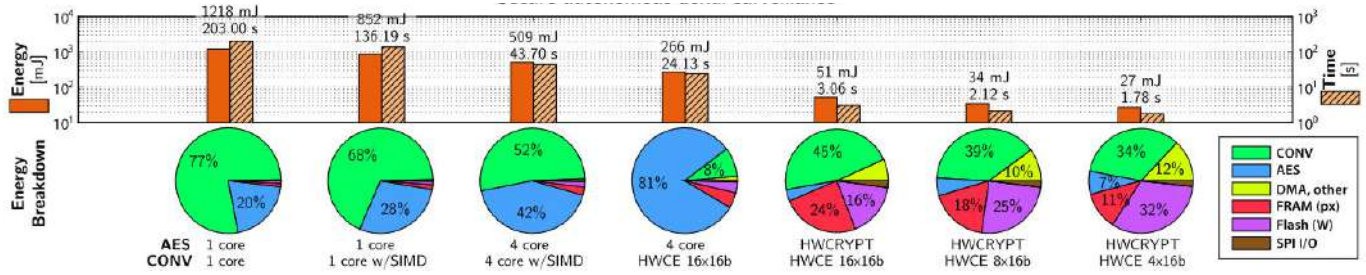


## Face Detection: 13x less energy

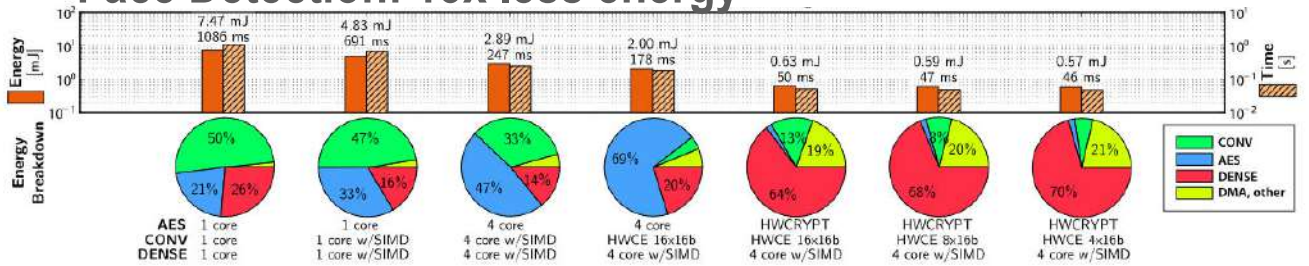


## EEG Based Seizure Detection: 2.1x less energy (Mostly SW)

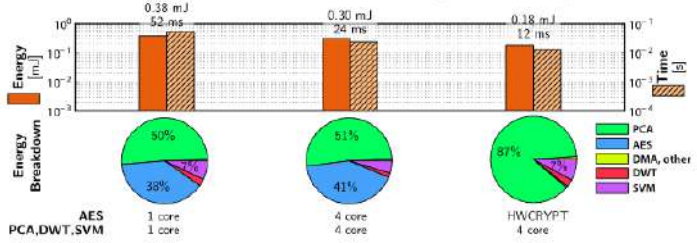


**Cluster good enough to run small end-to-end applications**
**I/O still problematic → Unbalance, Latency (not board ready…)**

# Fulmine: IEEE Best Paper Award

IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS–I: REGULAR PAPERS, VOL. 64, NO. 9, SEPTEMBER 2017 2481

## An IoT Endpoint System-on-Chip for Secure and Energy-Efficient Near-Sensor Analytics

Francesco Conti, Robert Schilling, Pasquale Davide Schiavone, Antonio Pullini, Davide Rossi, Frank Kağan Gürkaynak, Michael Muehlberghuber, Michael Gautschi, Igor Loi, Germain Haugou, Stefan Mangard, and Luca Benini

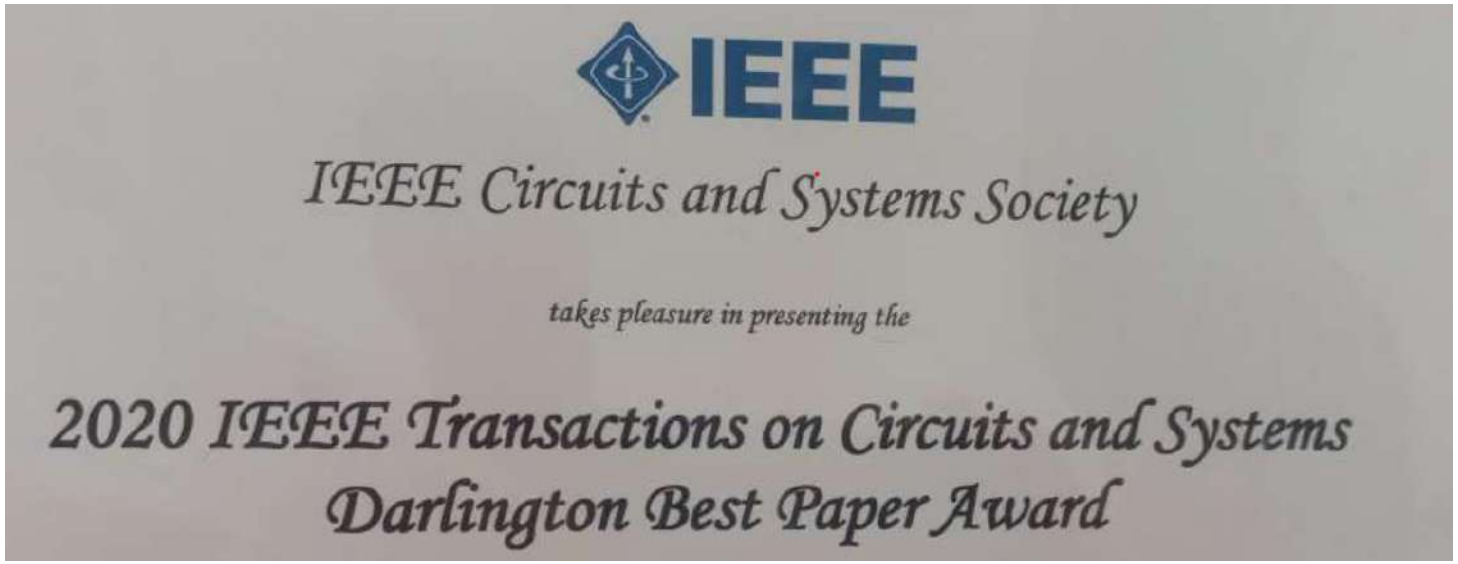

# Jan 2016: Transition to RISC-V Completed!

3-cycle ALU-OP, 4-cyle MEM-OP→ only IPC loss: LD-use, Branch

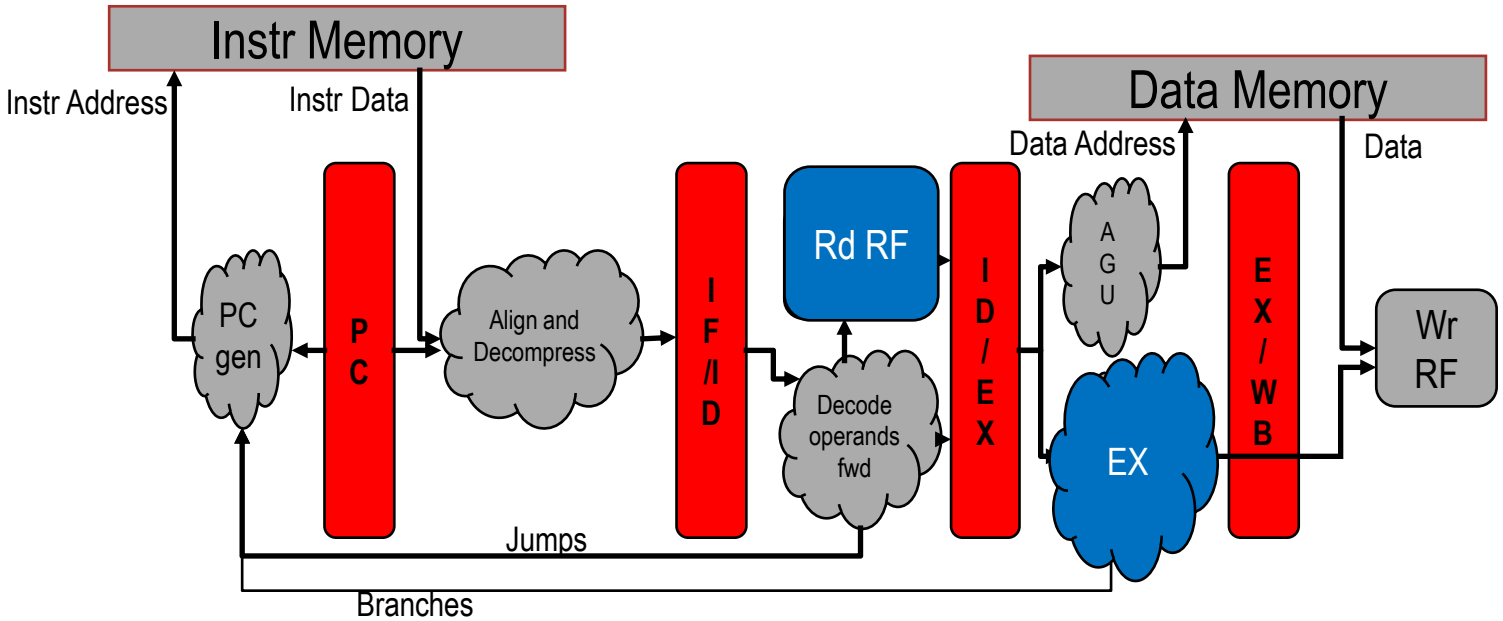


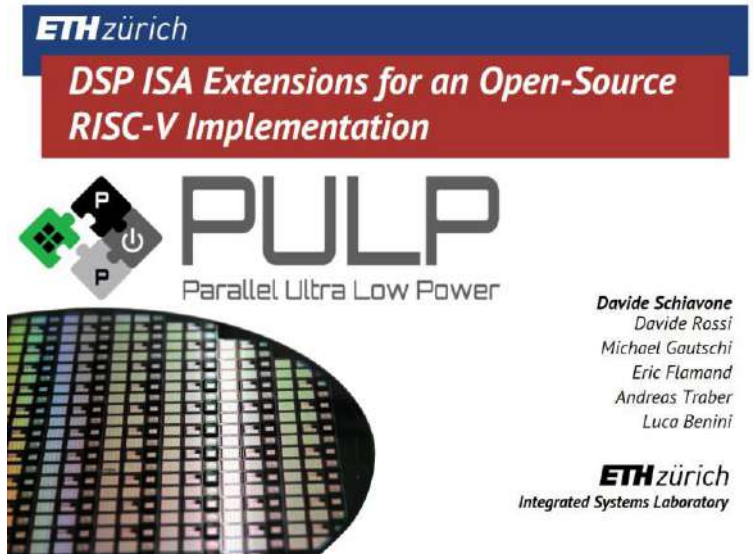


**V1** Baseline RISC

**V2** Extensions for Data Processing
Data motion (e.g. auto-increment)
Data processing (e.g. MAC)

**V3** Domain specific data processing
Narrow bitwidth
HW support for special arithmetic

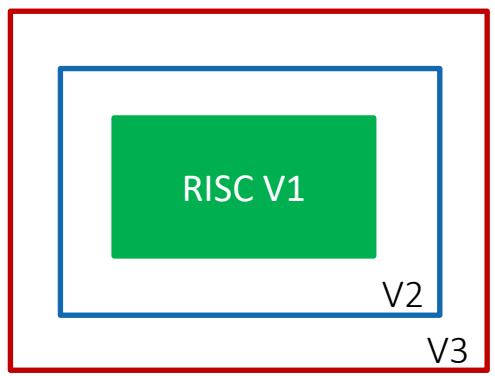

# First demonstration of RISC-V ISA Extensions

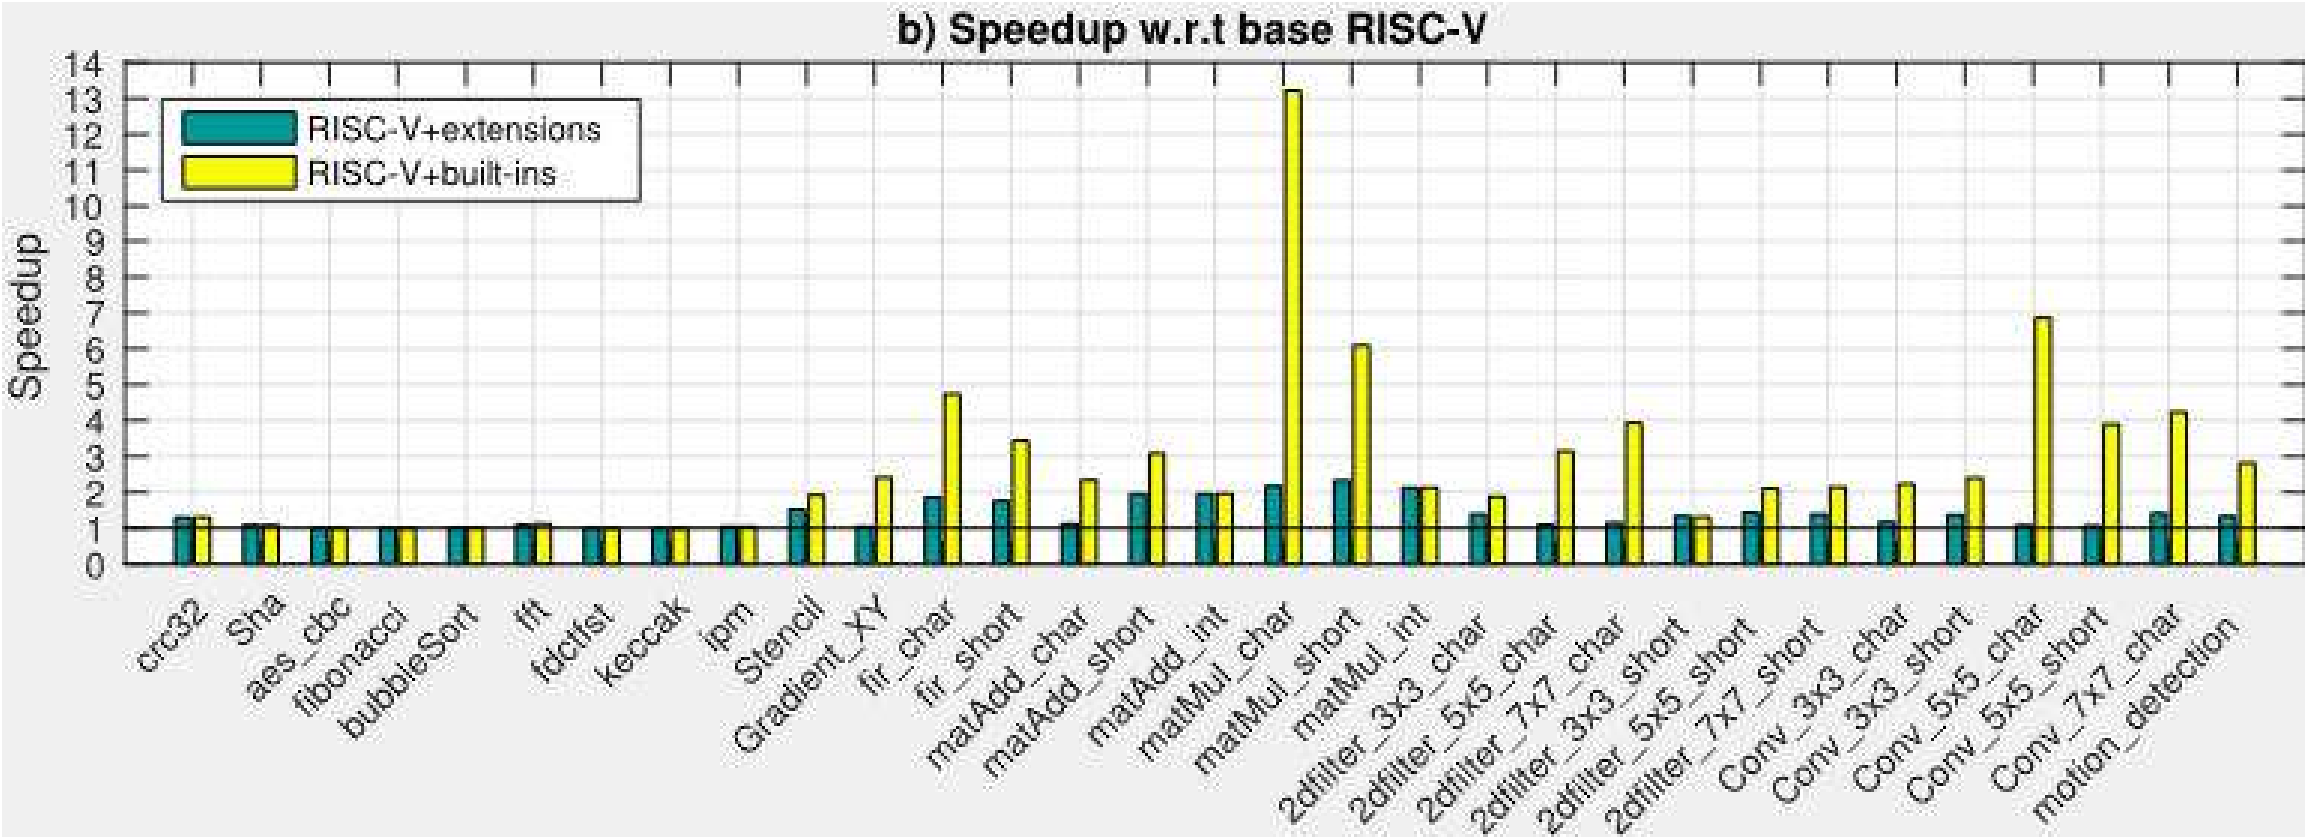

# IEEE TVLSI Best Paper Award

2700 IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS, VOL. 25, NO. 10, OCTOBER 2017

## Near-Threshold RISC-V Core With DSP Extensions for Scalable IoT Endpoint Devices

Michael Gautschi, *Student Member, IEEE*, Pasquale Davide Schiavone, *Student Member, IEEE*, Andreas Traber, Igor Loi, *Member, IEEE*, Antonio Pullini, *Student Member, IEEE*, Davide Rossi, *Member, IEEE*, Eric Flamand, Frank K. Gürkaynak, and Luca Benini, *Fellow, IEEE*

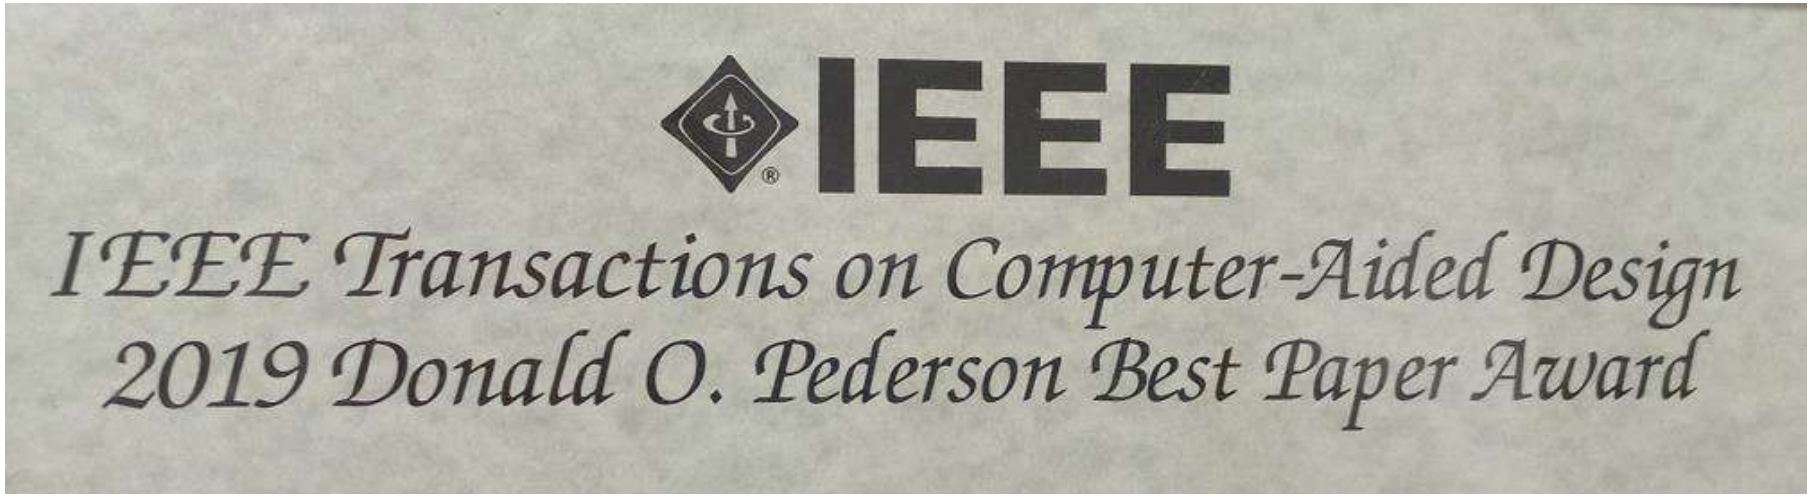

# Mr. Wolf SoC: First Board-Ready Chip (Dec 2018)

- **First SoC with Fabric controller: 32-bit RISC-V processor (Zero-RISKY)**
- **Rich set of peripherals + Autonomous IO DMA Subsystem (Improved)**
- **Parallel Programmable Accelerator:**
  - **First pulp chip with RI5CY**
  - 4 Shared Floating Point Units
  - AXI-4 DMA featuring 2D addressing capabilities
- **Advanced power management**
  - 2 Switchable Power Domains
  - 2 low-power FLLs (IO, SoC, Cluster)
  - On-chip DCDC and LDO (0.8V – 1.1V)

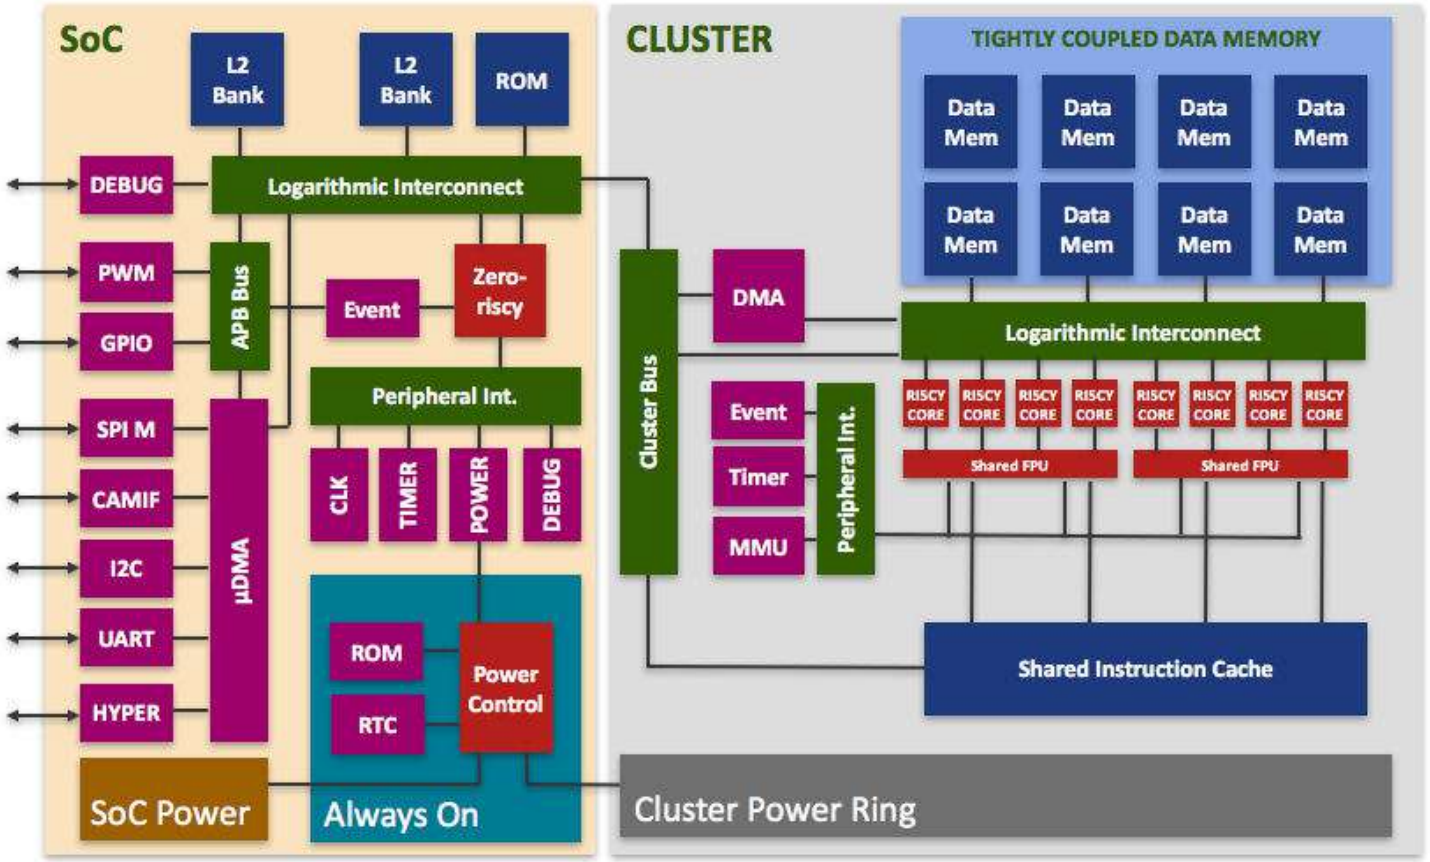

# Mr. Wolf @ ESSCIRC 2019

| | *PULP SoCs* | | |
|---|---|---|---|
| | PULPv2 [8] | FULMINE [9] | **Mr.Wolf (this work)** |
| Technology | FD-SOI 28nm flip-well | CMOS 65nm | **CMOS 40nm LP** |
| CPU | 32-bit OpenRISC | 32-bit OpenRISC | **32-bit RCVC32IMFX** |
| FPU | no | No | **2 shared** |
| # of cores | 4 | 4 | **1 + 8** |
| I$/D$/L2 | 1kBx4/ 48kB/ 64kB | 4K/ 64kB/ 192kB | **4K/ 64kB/ 512kB** |
| Power Management | clock gating | off-chip DC/DC | **DCDC + LDO + power gating** |
| Sleep Power | 200 µW | 120 µW | **72 – 108 µW state ret.** |
| IO Efficiency | Low | Mid | **High** |
| Voltage range (SRAMs) | 0.32V – 1.15V (0.45V – 1.15V) | 0.8V – 1.1V | **0.8V – 1.1V** |
| Max frequency | 825 MHz | 400 MHz | **450 MHz** |
| Best performance | 3.3 GOPS | 4.2 GOPS | **7 GOPS** |
| Best energy efficiency | 193 MOPS/mW @ 162 MOPS | 69 MOPS/mW @ 1.2 GOPS | **120 MOPS/mW @2.2 GOPS** |

Mr.Wolf: A 1 GFLOP/s Energy-Proportional Parallel Ultra Low Power SoC for IoT Edge Processing

Antonio Pullini*, Davide Rossi†, Igor Loi†, Alfio Di Mauro*, and Luca Benini*†
*Integrated Systems Laboratory, ETH Zürich, Gloriastr. 35, 8092 Zurich, Switzerland
†DEI, University of Bologna, Via Risorgimento 2, 40136 Bologna, Italy

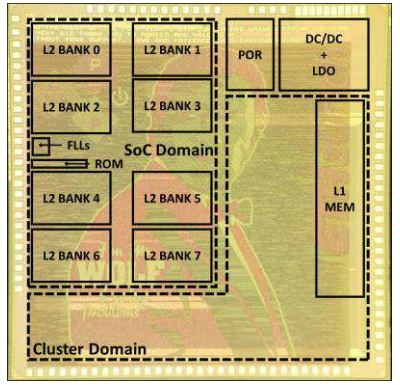



More on Simone Benatti's talk

More on Daniele Palossi's talk

# Vega: IoT Heterogeneous SoC

GREENWAVES TECHNOLOGIES

- Transprecision Floating-Point Units (16/32-bit float)[1]
- 4 MB non-volatile MRAM (Weight Storage)[2]
- Programmable Cognitive Wake-up Unit based on HD Computing

**ISSCC 2021 / SESSION 4 / PROCESSORS / 4.4**

**4.4 A 1.3TOPS/W @ 32GOPS Fully Integrated 10-Core SoC for IoT End-Nodes with 1.7µW Cognitive Wake-Up From MRAM-Based State-Retentive Sleep Mode**

Davide Rossi[1], Francesco Conti[1], Manuel Eggiman[2], Stefan Mach[2], Alfio Di Mauro[2], Marco Guermandi[1,3], Giuseppe Tagliavini[1], Antonio Pullini[2,3], Igor Loi[3], Jie Chen[1,3], Eric Flamand[2,3], Luca Benini[1,2]

[1]University of Bologna, Bologna, Italy
[2]ETH Zurich, Zurich, Switzerland
[3]Greenwaves Technologies, Grenoble, France

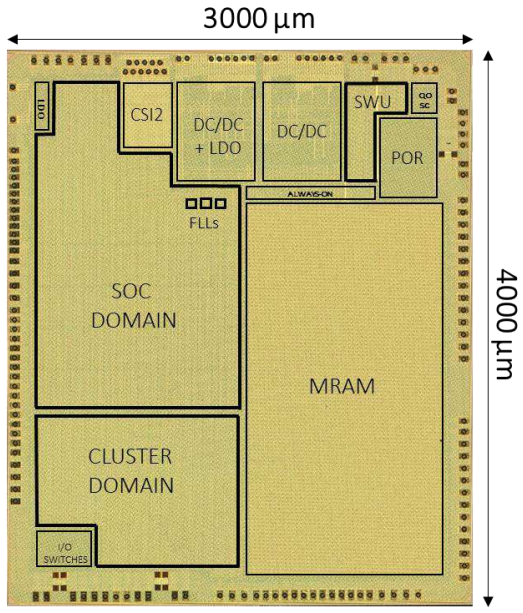


**Prototype implemented in GF 22FDX**

flip-well LVT & SLVT cells, 12mm2 for Full SoC

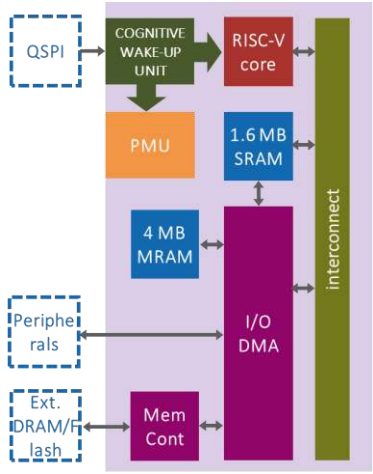


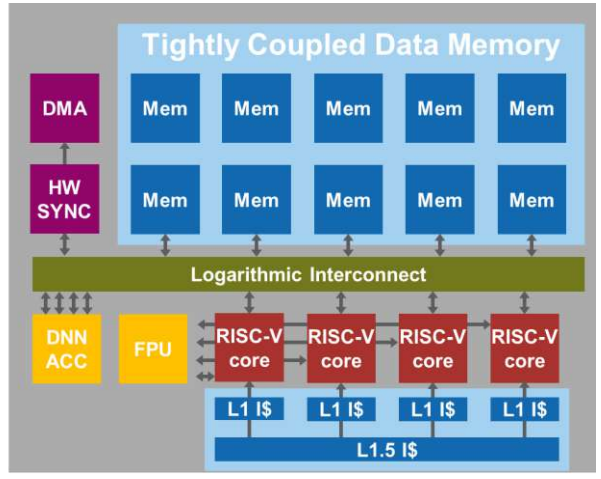


[1]More on Luca Bertaccini's talk
[2]More on Manuel Eggiman's talk

ETH zürich ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

# State-of-the-Art Comparison

| | SleepRunner [2] | Mr.Wolf [3] | SamurAI [4] | *VEGA* |
|---|---|---|---|---|
| Embedded NVM | - | - | SoA Int & FP Perf & Efficiency; Boost of 4.3x & 2.8x compared to Wolf | 4 MB MRAM |
| Wake-up Sources | WiC | GPIO, RT | | PIO, RTC, Cognitive |
| Best Int Perf.<br>Best.Int Eff.<br>@ Perf. | 31 MOPS (32b)<br>97 MOPS/mW (32b)<br>@ 18.6 MOPS (32b) | 12.1 GOPS<br>190 GOPS/W<br>@ 3.8 GOPS | 1.5 GOPS<br>230 GOPS/W<br>@110 MOPS | 15.6 GOPS<br>614 GOPS/W<br>@7.6 GOPS |
| Best FP Perf.<br>Best FP Eff.<br>@Perf | - | 1 GFLOPS<br>18 GFLOPS/W<br>@350 MFLOPS | - | 4 GFLOPS<br>158 GFLOPS/W<br>@ 2 GFLOPS |
| Best ML Perf.<br>Best ML Eff.<br>@Perf | - | - | 36 GOPS<br>1.3 TOPS/W<br>@ 2.8 GOPS | 32.2 GOPS<br>1.3 TOPS/W<br>@15.6 GOPS |

# First Fully Integrated MobileNetV2 on an IoT SoC

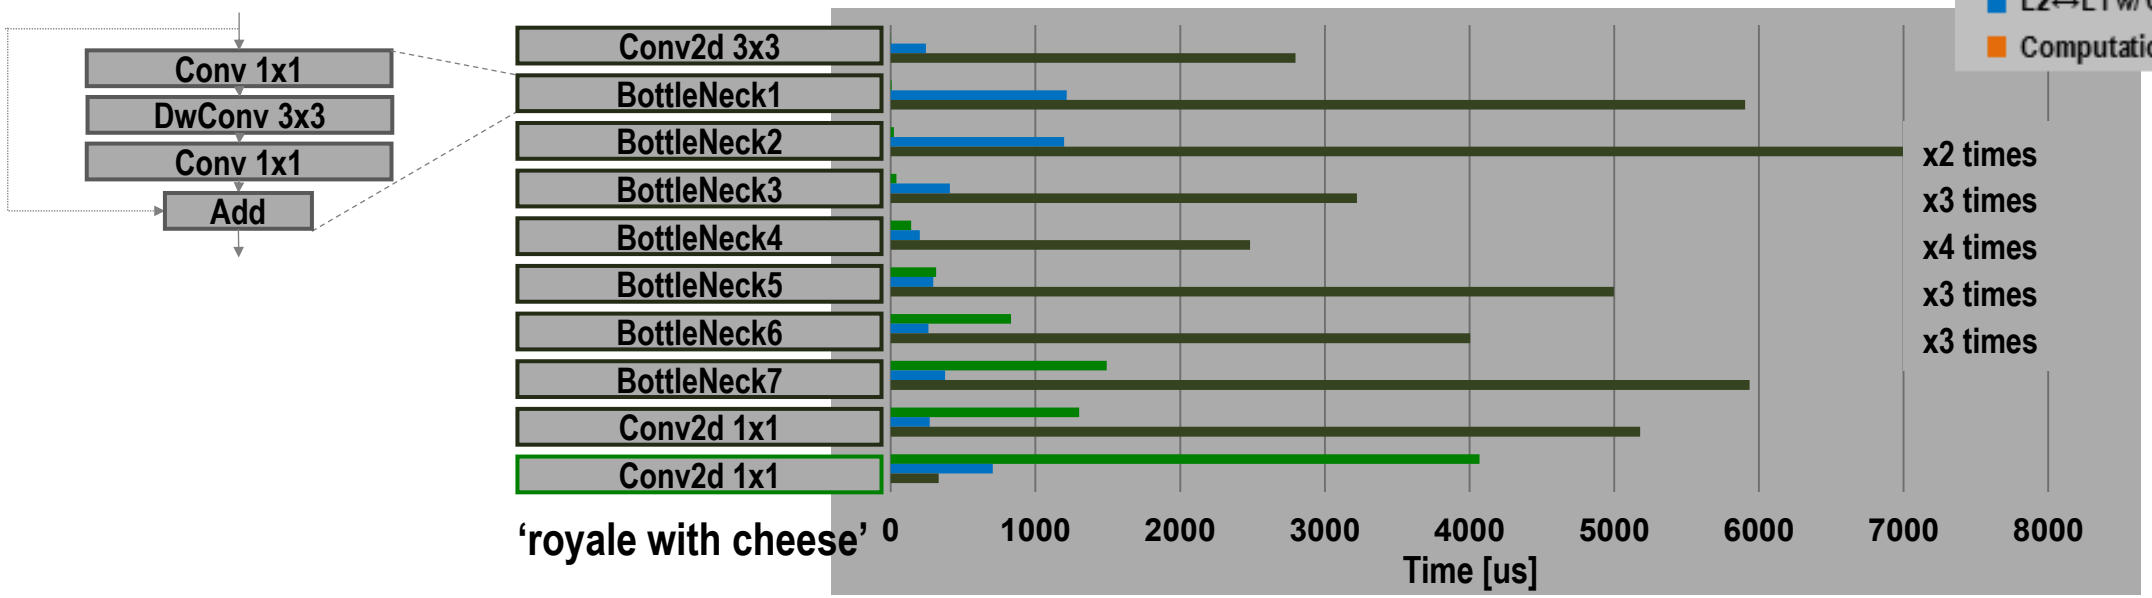


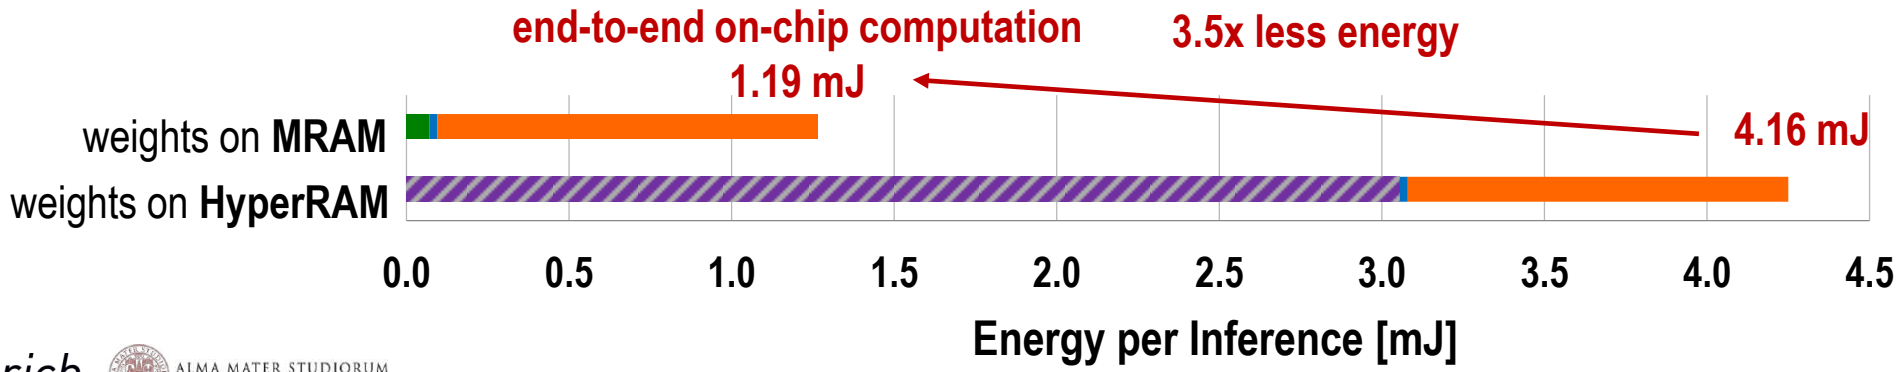

# MARSELLUS: AI-IoT Heterogeneous SoC

- XpulpNN + M&L Extensions
- 2-8b Reconfigurable Binary Engine for 3x3, 1x1 DNN kernels
- Adaptive Body Biasing with on-the-fly control

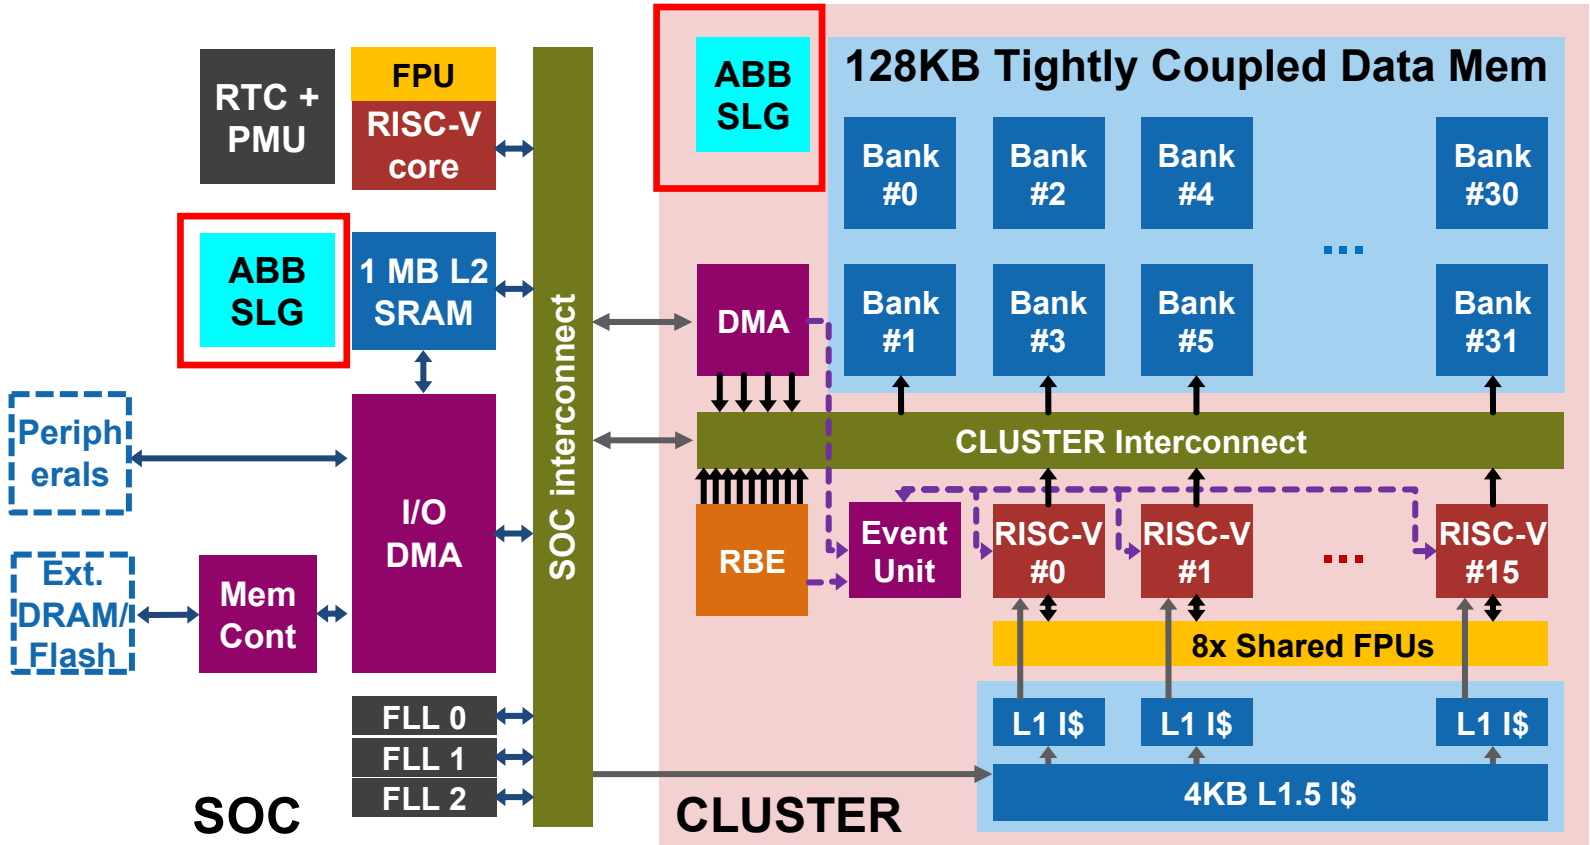

# MARSELLUS: Presented at ISSCC 2023

- XpulpNN + M&L Extensions
- 2-8b Reconfigurable Binary Engine for 3x3, 1x1 DNN kernels
- Adaptive Body Biasing with on-the-fly control

22.1 A 12.4TOPS/W @ 136GOPS AI-IoT System-on-Chip with 16 RISC-V, 2-to-8b Precision-Scalable DNN Acceleration and 30%-Boost Adaptive Body Biasing<br>

Francesco Conti[1], Davide Rossi[1], Gianna Paulin[2], Angelo Garofalo[1], Alfio Di Mauro[2], Georg Ruetishauer[2], Gianmarco Ottavi[1], Manuel Eggimann[2], Hayate Okuhara[1], Vincent Huard[3], Olivier Montfort[3], Lionel Jure[3], Nils Exibard[3], Pascal Gouedo[3], Mathieu Louvat[3], Emmanuel Botte[3], Luca Benini[1,2]

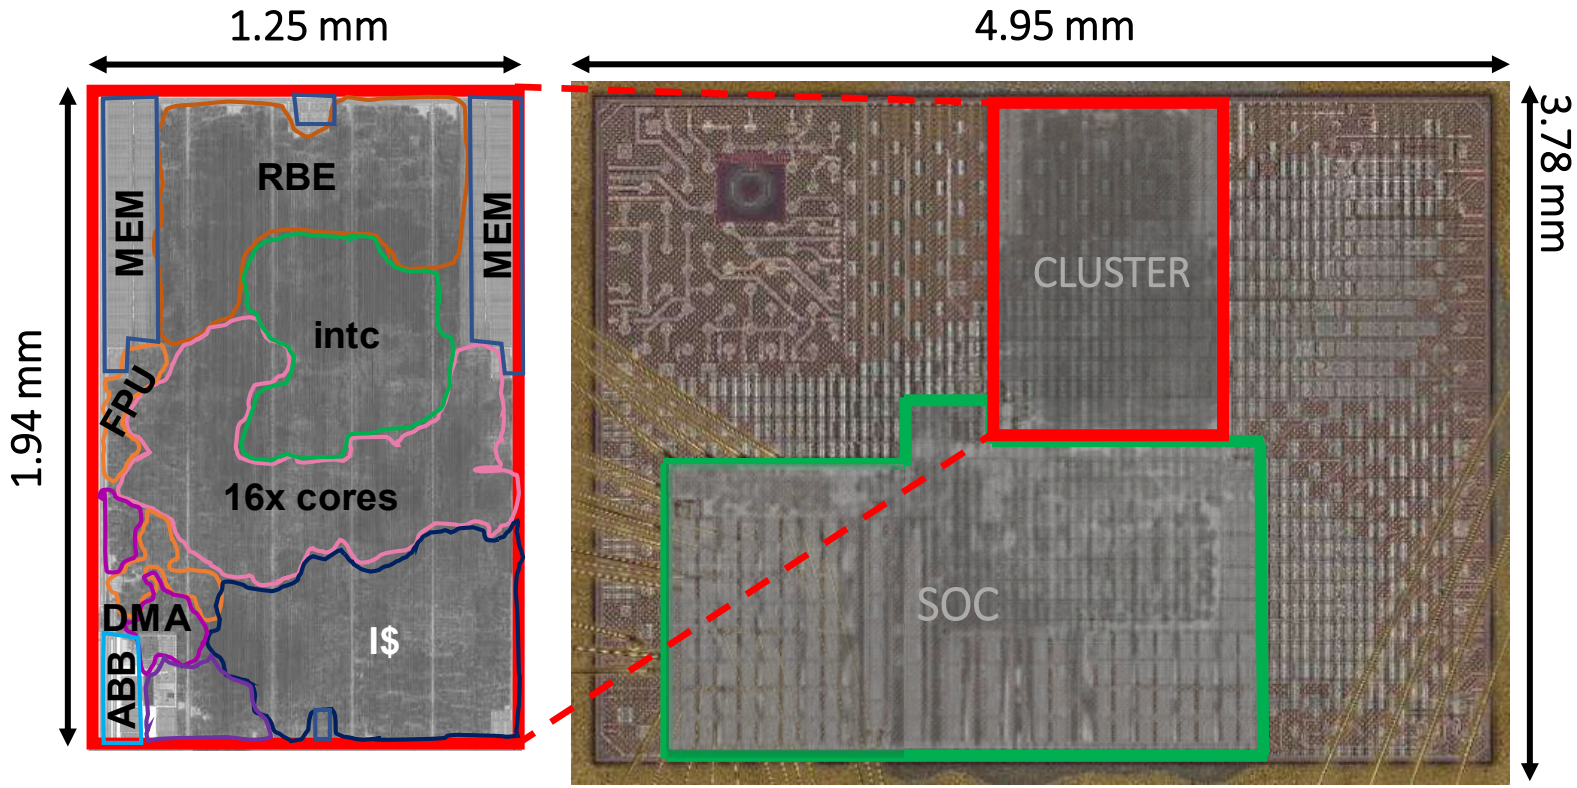

# State-of-the-Art Comparison (SW Programmable)

| | VEGA [1] ISSCC'21 | SAMURAI [2] VLSIC'20 | DIANA [3] ISSCC'22 | *MARSELLUS* |
|---|---|---|---|---|
| **Technology** | 22nm FDX | 28nm FD-SOI | 22nm FDX + AIMC | 22nm FDX |
| **Die Area** | 10 $mm^2$ | 4.5 $mm^2$ | 10.24 $mm^2$ | 18.7 $mm^2$ (cluster 1.9 $mm^2$) |
| **Cores** | 10x RV32IMCFXpulp + HWCE | 1x RV32IMCFXpulp + Digital Accel. | 1x RV32CIMFXpulp + Digital Accel. + AIMC | 1x RV32IMCFXpulp + 16x RV32IMCFXpulpnn + RBE |
| **Max Frequency** | 450 MHz | 350 MHz | 320 MHz | 420 MHz |
| **Power range** | 1.7 uW - 49.4 mW | 6.4 uW - 96 mW | 10-129 mW (digital) | 12.8 mW - 123 mW |
| **Best SW (INT) Perf** | 15.6 GOPS (8 RISC-V) | 1.5 GOPS (1 | | **90 GOPS (16 RISC-V M&L 2x2b, 0.8V+ABB)** |
| **Best SW (INT) Eff** | 614 GOPS/W @ 7.6 GOPS (8 RISC-V) | 230 @ 1 (1 | | **1.66 TOPS/W @ 19 GOPS (16 RISC-V M&L 2x2b)** |
| **Best HW-Accel Perf** | 32.2 GOPS (HWCE) | 36 GOPS (Digital Accel.) | 180 GOPS (Digital), **29.5 TOPS (AIMC)** | **637 GOPS (RBE 2x2b, 0.8V+ABB)** |
| **Best HW-Accel Eff** | 1.3 TOPS/W @ 15.6 GOPS (HWCE) | 1.3 TOPS/W @ 2.8 GOPS (Digital Accel.) | 4.1 TOPS/W (Digital), **600 TOPS/W (AIMC)** | **12.4 TOPS/W @ 136 GOPS (RBE 2x2b)** |

**Machine Learning SW Perf & Efficiency boost of 5.8x & 2.7x compared to VEGA**

# State-of-the-Art Comparison (HW Accelerated)

| | VEGA [1] ISSCC'21 | SAMURAI [2] VLSIC'20 | DIANA [3] ISSCC'22 | *MARSELLUS (this work)* |
|---|---|---|---|---|
| **Technology** | 22nm FDX | 28nm FD-SOI | 22nm FDX + AIMC | 22nm FDX |
| **Die Area** | 10 $mm^2$ | 4.5 $mm^2$ | 10.24 $mm^2$ | 18.7 $mm^2$ (cluster 1.9 $mm^2$) |
| **Cores** | 10x RV32IMCFXpulp + HWCE | 1x RV32IMCFXpulp + Digital Accel. | 1x RV32CIMFXpulp + Digital Accel. + AIMC | 1x RV32IMCFXpulp + 16x RV32IMCFXpulpnn + RBE |
| **Max Frequency** | 450 MHz | 350 MHz | 320 MHz | 420 MHz |
| **Power range** | 1.7 uW - 49.4 mW | 6.4 uW - 96 mW | 10-129 mW (digital) | 12.8 mW - 123 mW |
| **Best SW (INT) Perf** | 15.6 GOPS (8 RISC-V) | 1.5 GOPS (1 RISC-V) | - | **90 GOPS (16 RISC-V M&L 2x2b, 0.8V+ABB)** |
| **Best SW (INT) Eff** | 614 GOPS/W @ 7.6 GOPS (8 RISC-V) | 230 GOPS/W @ 110 MOPS (1 RISC-V) | - | **1.66 TOPS/W @ 19 GOPS (16 RISC-V M&L 2x2b)** |
| **Best HW-Accel Perf** | 32.2 GOPS (HWCE) | 36 GOPS (Dig | 180 GOPS (Digital), AIMC) | **637 GOPS (RBE 2x2b, 0.8V+ABB)** |
| **Best HW-Accel Eff** | 1.3 TOPS/W @ 15.6 GOPS (HWCE) | 1.3 @ 2 (Digital Accel.) | /W 600 TOPS/W (AIMC) | **12.4 TOPS/W @ 136 GOPS (RBE 2x2b)** |

DNN Acceleration
Perf & Efficiency boost
>10x w.r.t. VEGA



More on Francesco Conti's talk

# Shaheen: Heterogeneous Application Class SoC

- **Secure Heterogeneous Application Processor**
- **Host Subsystem**
  - CVA6
  - H-Extensions →BAO, Sel4
  - Timing channel attack protection
- **PULP Cluster with Mixed-Precision Extension**
- **HyperRAM Memory Controller**
  - Up to 512 Mbit
  - Up to 1.6 Gbit/s

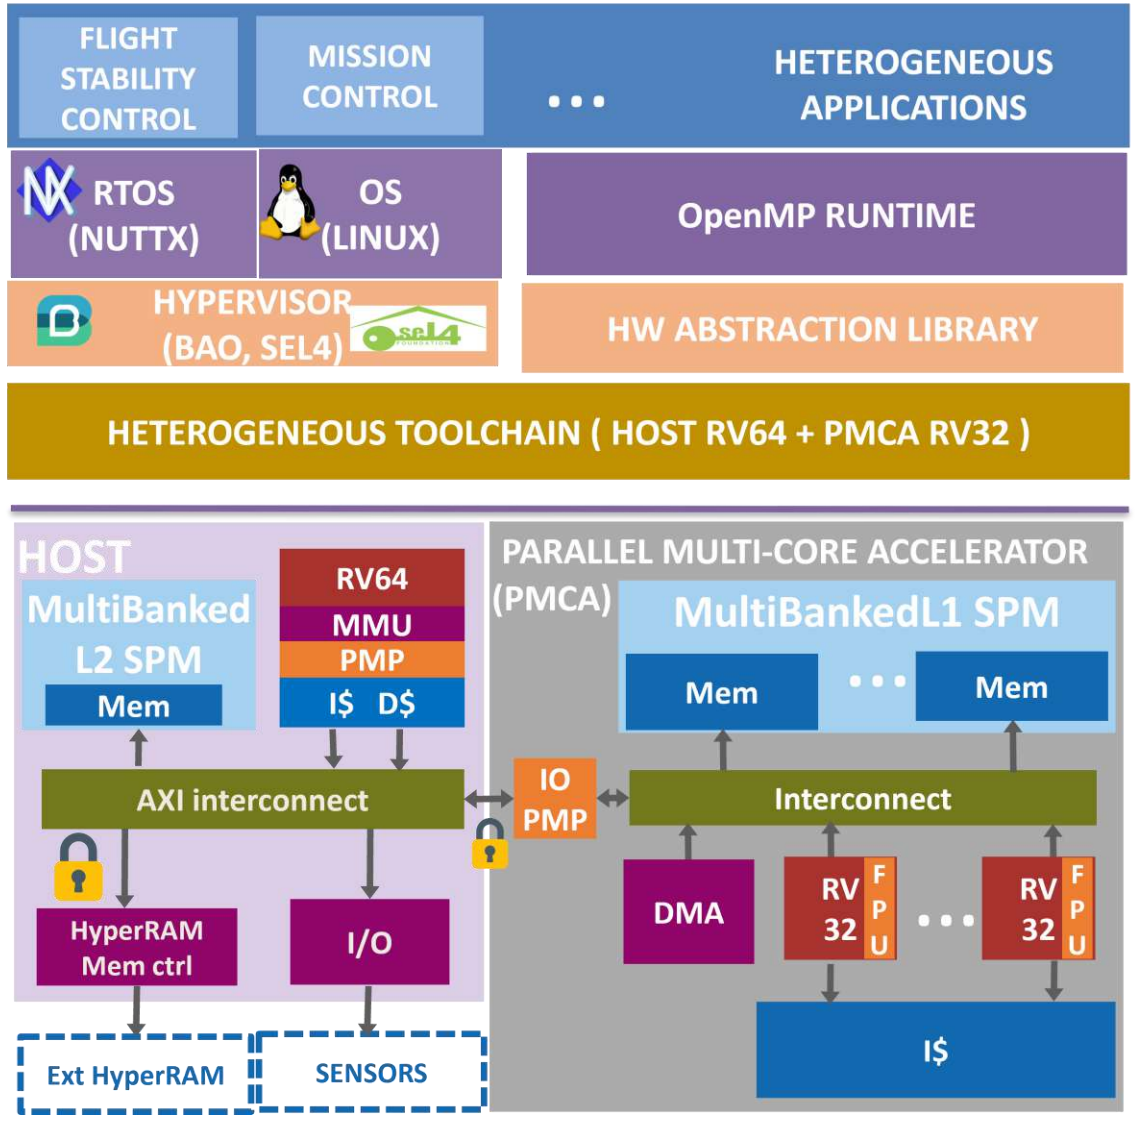


More on Angelo Garofalo's talk

# Shaheen to be presented @ Hot Chips 2023

## Shaheen: An Open, Secure, and Scalable RV64 SoC for Autonomous Nano-UAVs

L. Valente*, A. Veeran¶, M. Sinigaglia*, Y. Tortorella*, A. Nadalini*, N. Wistoff†, B. Sá‡, A. Garofalo*, R. Psiakis**, M. Tolba¶, A. Kulmala**, N. Limaye, O. Sinanoglu§, S. Pinto‡, D. Palossi†‖, L. Benini*†, B. Mohammad¶, D. Rossi*
* University of Bologna, Italy, ¶ Khalifa University, UAE, † ETH Zürich, Switzerland, ‡ Universidade do Minho, Portugal,
** Technology Innovation Institute, UAE § New York University Abu Dhabi, UAE, ‖ IDSIA USI-SUPSI, Switzerland

HOT CHIPS

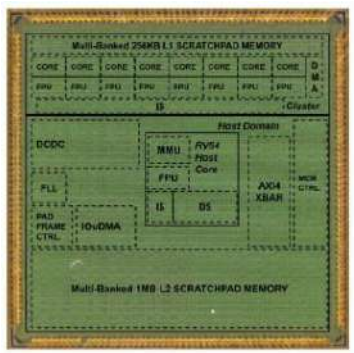

Single SoC mission computer for nano-drones

Flight computer for standard drones

Universidade do Minho

جامعة خليفة Khalifa University

TII Technology Innovation Institute

USI/SUPSI IDSIA

جامعة نيويورك أبوظبي NYU ABU DHABI

ETH zürich ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

# PULP Master Student Chips @ ESSCIRC 21, 22, ISCAS 2023

## *Dustin,* ESSCIRC 2021

A. Garofalo et al.

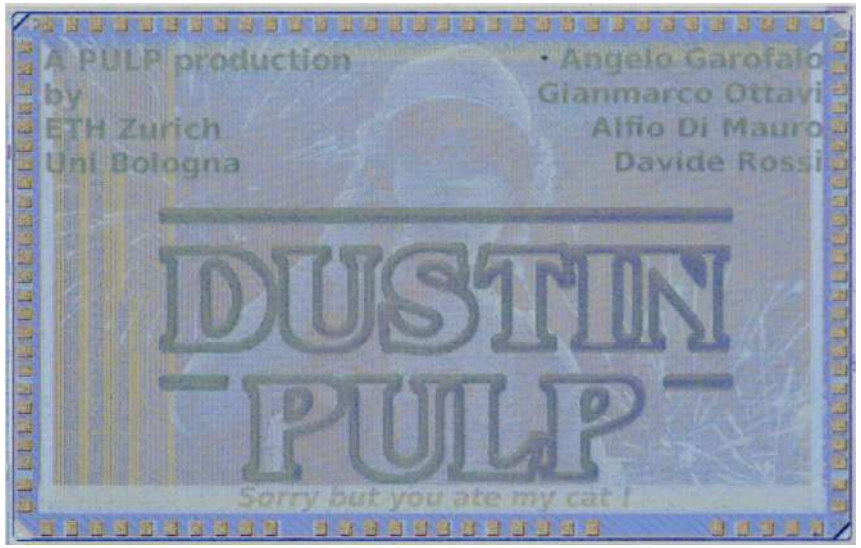


16-Cores Vector Lockstep Execution Mode RISC/V Cluster for Inference of Quantized Neural Networks in 65 nm CMOS.

## *Darkside, ESSCIRC 2022*

A. Garofalo, et al.


Specialized Heterogeneous RISC-V Cluster for Inference and Training of Deep Neural Networks in 65nm CMOS.

## *Echoes,* ISCAS 2023

M. Sinigaglia, et al.


Specialized Heterogeneous SoC for Audio and Frequency-Domain Processing in 65nm CMOS.


# PULP

Parallel Ultra Low Power

# Questions?

 http://pulp-platform.org  @pulp_platform